

AMENDMENTS TO SPECIFICATION

Please amend the BACKGROUND OF THE INVENTION section as follows.

[0004] Enterprises have are looking at ways of reducing costs and increasing efficiencies of their data processing systems that serve the needs of many applicationssystem. A typical enterprise data processing system allocates individual resources for each of the enterprise's applications. Application workEnough resources are acquired for each application to handle the estimated peak load of the application. Each application has different load characteristics are not; some applications are busy during the day; some others during the night; some reports are run once a week and some others once a month. Enough resources are acquired to ensure that the peak work load of each application may be handled. As a result, when the workload of an applicationthere is not at its peak, a lot of resource allocated to an applicationcapacity that is left unutilized. Grid computing enables the utilization or elimination of this unutilized capacity. In fact, Grid computing is poised to drastically change the economics of computing.

[0005] Based on the foregoing, there is need to provide a mechanism that meet vary workload requirements of applications in a way

A grid is a collection of commodity computing elements that provide processing and some degree of shared storage; the resources of a grid are allocated dynamically to meet the computational needs and priorities of its clients. An example of a grid is a rack of server blades. Each server blade is an inclusive computer system, with processor, memory, network connections, and associated electronics on a single motherboard. Typically, server blades do not include onboard storage (other than volatile memory), and they share storage units (e.g. shared disks) along with a power supply, cooling system, and cabling within a rack.

Grid computing can dramatically lower the cost of computing, extend the availability of computing resources, and deliver higher productivity and higher quality. The basic idea of Grid computing is the notion of computing as a utility, analogous to the electric power grid or the telephone network. A client of the Grid does not care where its data is or where the computation is performed. All a client wants is to have computation done and have the information delivered to the client when it wants.

At a high level, the central idea of Grid computing is computing as a utility. A client of a Grid should not have to care where its data resides, or what computer element processes a request. The client need only request information or computation and have it delivered – as much as needed and whenever needed. This is analogous to the way electric utilities work; a customer does not know where the generator is, or how the electric grid is wired. The customer just asks for electricity and gets it. The goal is to make computing a utility -- a ubiquitous commodity. Hence it has the name, the Grid.

This view of Grid computing as a utility is, of course, a client side view. From the server side, or behind the scenes, the Grid is about resource allocation, information sharing, and high availability. Resource allocation ensures that all those that reduces the need to allocate an overcapacity of resources that is left unused during non-peak load times or request resources are getting what they need. Resources are not standing idle while requests are left unserviced. Information sharing makes sure that the information clients and applications need is available where and when it is needed. High availability ensures that all the data and computation must always be there -- just as a utility company must always provide electric power.

### Grid Computing for Databases

One area of computer technology that can benefit from Grid computing is database technology. A grid can support multiple databases and dynamically allocate resources as needed

to support the load on each database. As the load for a database increases, more resources are allocated for that database. For example, on an enterprise grid, a database is being serviced by one database server running on one server blade on the grid. The number of users requesting data from a database increases. In response to this increase in the demand for the database, another database server is provisioned on one or more other server blades.

#### Provisioning for Database Grid

The term provisioning refers to providing and configuring the computational resources and data needed to provide a service. Provisioning a database server includes configuring a computer element to run the database server and configuring the database server to manage a database. Provisioning a database includes moving a database to where a database server may access it and configuring the database server to manage the database.

The process of provisioning data or a database is referred to herein as data provisioning. Provisioning a database in a grid may require cloning all or part of the database, followed by provisioning a new database server to manage the clone or incorporating the clone into another database already being managed by an already running database server.

Data provisioning of a database can also involve the bulk transfer of data between file systems and/or databases. Unfortunately, techniques for bulk transfer of data that are used for database provisioning entail manual intervention and therefore cannot be used to effectively provision data automatically and dynamically as is required for grid computing.

An example of an approach for data provisioning that uses a technique for bulk transfer of data is the transportable tablespace approach. A tablespace is a collection of storage containers (e.g. files) used to store data for database objects (e.g. relational tables). Under this approach, tablespaces are exported from a “source database” and imported into a “target database”. This capability allows the files of a tablespace to be copied using operating system utilities for

copying files, which run much faster than the other techniques for bulk transfer of data between databases. Such other techniques involve executing queries and insert statements.

To transport a tablespace, a human database administrator ("DBA") performs manual steps. First, the tablespace must be imported into the target tablespace by attaching the tablespace. With respect to a tablespace, database, and database server, the term "attach" refers to configuring a database and/or database server so that the database objects in the tablespace are incorporated within the database and the tablespace is used to store data for the database. Configuring a database to attach a tablespace involves modifying database metadata so that it defines the tablespace and database objects as part of the database. The database metadata may be altered using a variety of techniques involving manual steps performed by a DBA. The DBA can run utilities available on the source database system that may be executed to export the metadata into a "metadata dump file", and run utilities on the target database system to construct metadata from the metadata dump file. Alternately, metadata can be included with the data being transported in the tablespace, and the target database would reconstruct the metadata from the metadata included in the tablespace. The DBA can also manually reconstruct the metadata on the target database system.

A tablespace may be transported to a database by creating a separate copy of the tablespace from the original source database and attaching it to the target database. While the copy is being made, operations on the tablespace are restricted to read-only operations. The DBA sends commands to instruct a database server managing the database to restrict database operations performed on the tablespace to read-only operations. Once the copy is complete, the DBA may send commands to instruct the database server that modification operations can be performed.

The term "copy," as used herein, refers to both the source data and a duplicate of the source data. For example, a copy of a source file may be the source file itself, or another file that is a duplicate that can be generated using, for example, readily available copy utilities, such as operating system utilities for creating copies of files.

The copy of a tablespace transported may also be detached from a database. With respect to a particular tablespace and database and database server, the term detach refers to configuring a database and/or database server so that a tablespace is no longer used to store data for the database. Configuring a database to detach a tablespace includes altering database metadata in the source database system, by, for example, removing metadata defining the tablespace as part of the source database system, or setting a flag to indicate that the tablespace is no longer used. This step is performed by the DBA by running utilities or by manually editing the source database metadata.

The tablespaces for the source database are stored in a "source directory" of a file system and the tablespaces of the target database system are stored in a "target directory" of a file system. The source and target directories may be within the same file system or a different file system on a different computer system. In any case, the DBA needs to transfer the tablespace using operating system utilities. This requires that the DBA have an operating system account on the computer system of the target directory. The DBA logs onto the computer system and runs a utility to transfer the tables from the source directory to the target directory. If the target directory is in another file system of another computer system, the DBA can use FTP to transfer the file (i.e. use a utility that follows the File Transfer Protocol). To use FTP, the DBA needs an operating system account on the computer system of the target directory to log onto the computer system and transfer the tablespace files.

As demonstrated above, conventional bulk transfer of data techniques for provisioning a database require manual intervention on the part of a human DBA. Because Grid computing requires that data provisioning be performed automatically and dynamically, database provisioning that requires the bulk transfer of data is not amenable for Grid computing. Clearly, there is need for automated bulk transferring of databases that is suitable for dynamic data provisioning within a grid.

[0006] —— The approaches described in this section are approaches that could be pursued, but not necessarily approaches that have been previously conceived or pursued. Therefore, unless otherwise indicated, it should not be assumed that any of the approaches described in this section qualify as prior art merely by virtue of their inclusion in this section.